

Compact preference representation for Boolean games

Elise Bonzon, Marie-Christine Lagasquie-Schiex, and Jérôme Lang

IRIT, UPS, F-31062 Toulouse Cédex 9, France

{bonzon,lagasq,lang}@irit.fr

Abstract

Boolean games, introduced by [15, 14], allow for expressing compactly two-players zero-sum static games with binary preferences: an agent's strategy consists of a truth assignment of the propositional variables she controls, and a player's preferences is expressed by a plain propositional formula. These restrictions (two-players, zero-sum, binary preferences) strongly limit the expressivity of the framework. While the first two can be easily encompassed by defining the agents' preferences as an arbitrary n -uple of propositional formulas, relaxing the last one needs Boolean games to be coupled with a propositional language for compact preference representation. In this paper, we consider generalized Boolean games where players' preferences are expressed within two of these languages: prioritized goals and propositionalized CP-nets.

1 Introduction

The framework of Boolean games, introduced by [15, 14], allows for expressing compactly two-players zero-sum static games with binary preferences: an agent's strategy consists of a truth assignment of the propositional variables she controls, and a player's preferences is expressed by a plain propositional formula. Arguably, these three restrictions (two-players, zero-sum, binary preferences) strongly limit the expressivity of the framework. The first two can be easily encompassed by defining the agents' preferences as an arbitrary n -uple of propositional formulas (see [3], who addresses complexity issues for these binary n -players Boolean games). In this paper we focus on the third one, which needs considerably more work to be dealt with. The starting point of our paper is that whereas a single propositional formula (goal) ϕ cannot express more than a binary preference relation on interpretations (models of ϕ are strictly better than models of $\neg\phi$), expressing arbitrary (non-binary) preferences within a propositional framework is possible, making use of a *propositional language for compact preference representation*. The study of such languages has been a very active issue for a few years in the AI community. Several classes of languages based on propositional logic have been proposed and studied (see for instance [16, 8] for an overview of these languages).

A first question has to be addressed before going further: should agents' preferences be expressed in a numerical way or in an ordinal way? This depends a lot on the notions we want to deal with. While some notions (such as pure Nash equilibria and dominated strategies) can be defined in a purely ordinal setting, other ones (such as mixed strategy Nash equilibria) need quantitative (real-valued) preferences. Here we choose to stick to ordinal settings (we leave numerical preferences in Boolean games for further work – see Section 5), and we successively integrate Boolean games with two of these languages: first, *prioritized goals*, and then (*propositionalized*) *CP-nets*.

In Section 2, some background is given and we define n -players, non zero-sum) Boolean games with binary preferences. Boolean games are then enriched with prioritized goals in Section 3, and with propositionalized CP-nets in Section 4. Section 5 addresses related work and further issues.

Copyright ©2006 : Springer-Verlag. This paper appears in Lecture Note in Computer Science in a volume devoted to PRICAI'06. All rights reserved.

2 n -players Boolean games

Let $V = \{a, b, \dots\}$ be a finite set of propositional variables and L_V be the propositional language built from V and the usual connectives as well as the Boolean constants \top (*true*) and \perp (*false*). Formulas of L_V are denoted by ϕ, ψ , etc. A *literal* is a formula of the form x or of the form $\neg x$, where $x \in V$. A *term* is a consistent conjunction of literals.

2^V is the set of the interpretations for V , with the usual meaning that an interpretation M gives the value *true* to a variable x if $x \in M$, and the value *false* otherwise. \models denotes classical logical consequence. Let $X \subseteq V$. 2^X is the set of X -interpretations. A *partial interpretation* (for V) is an X -interpretation for some $X \subseteq V$. Partial interpretations are denoted by listing all variables of X , with a $\bar{}$ symbol when the variable is set to false: for instance, let $X = \{a, b, d\}$, then the X -interpretation $M = \{a, d\}$ is denoted $a\bar{b}d$. If $\{V_1, \dots, V_p\}$ is a partition of V and $\{M_1, \dots, M_p\}$ are partial interpretations, where $M_i \in 2^{V_i}$, (M_1, \dots, M_p) denotes the interpretation $M_1 \cup \dots \cup M_p$.

Given a set of propositional variables V , a Boolean game on V [15, 14] is a zero-sum game with *two players* (1 and 2), where the actions available to each player consist in assigning a truth value to each variable in a given subset of V . The utility functions of the two players are represented by a propositional formula ϕ formed upon the variables in V and called *Boolean form* of the game¹. ϕ represents the goal of Player 1: her payoff is 1 when ϕ is satisfied, and 0 otherwise. Since the game is zero-sum², the goal of Player 2 is $\neg\phi$. This simple framework can be extended in a straightforward way to non zero-sum n -players games (see [3], especially for complexity issues): each player i has a goal ϕ_i (a formula of L_V). Her payoff is 1 when ϕ_i is satisfied, and 0 otherwise.

Definition 1 A n -players Boolean game is a 4-uple (A, V, π, Φ) , where $A = \{1, 2, \dots, n\}$ is a set of players, V is a set of propositional variables $\pi : A \mapsto V$ is a control assignment function and $\Phi = \langle \phi_1, \dots, \phi_n \rangle$ is a collection of formulas of L_V .

The control assignment function π associates every player with the variables that she controls. For the sake of notation, the set of all the variables controlled by i is written π_i instead of $\pi(i)$. We require that each variable be controlled by one and only one agent, i.e., $\{\pi_1, \dots, \pi_n\}$ forms a partition of V . The original definition by [15, 14] is a special case of this more general framework, obtained by letting $n = 2$ and $\phi_2 = \neg\phi_1$.

Definition 2 Let $G = (A, V, \pi, \Phi)$. A **strategy** s_i for a player i is a π_i -interpretation. A **strategy profile** S for G is an n -uple $S = (s_1, s_2, \dots, s_n)$ where for all i , $s_i \in 2^{\pi_i}$.

In other words, a strategy for i is a truth assignment for all the variables i controls. Remark that since $\{\pi_1, \dots, \pi_n\}$ forms a partition of V , a strategy profile S is an interpretation for V , i.e., $S \in 2^V$. Ω denotes the set of all strategy profiles for G .

The following notations are usual in game theory. Let $G = (A, V, \pi, \Phi)$, $S = (s_1, \dots, s_n)$, $S' = (s'_1, \dots, s'_n)$ be two strategy profiles for G . s_{-i} denotes the projection of S on $A \setminus \{i\}$: $s_{-i} = (s_1, s_2, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$. Similarly, π_{-i} denotes the set of the variables controlled by all players except i : $\pi_{-i} = V \setminus \pi_i$. Finally, (s_{-i}, s'_i) denotes the strategy profile obtained from S by replacing s_i with s'_i without changing the other strategies: $(s_{-i}, s'_i) = (s_1, s_2, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$.

Example 1 We consider here a Boolean n -players version of the well-known prisoners' dilemma. n prisoners (denoted by $1, \dots, n$) are kept in separate cells. The same proposal is made to each of them: "Either you cover your accomplices (C_i , $i = 1, \dots, n$) or you denounce them ($\neg C_i$, $i = 1, \dots, n$). Denouncing makes you freed while your partners will be sent to prison (except those who denounced you as well; these ones will be freed as well). But if none of you chooses to denounce, everyone will be freed."³ This can be expressed much compactly by the following n -players Boolean game $G = (A, V, \pi, \Phi)$:

¹The original definition in [15, 14] is inductive: a Boolean game consists of a finite dynamic game. We use here the equivalent, simpler definition of [11], who showed that this tree-like construction is unnecessary.

²Stricto sensu, the obtained games are not zero-sum, but constant-sum (the sum of utilities being 1) – the difference is irrelevant and we use the terminology "zero-sum" nevertheless.

³The case where everyone will be freed if everyone denounces the others is a side effect of our simplification of the prisoners' dilemma.

strategy of 3: C_3			strategy of 3: \overline{C}_3		
1 \ 2	C_2	\overline{C}_2	1 \ 2	C_2	\overline{C}_2
C_1	(1, 1, 1)	(0, 1, 0)	C_1	(0, 0, 1)	(0, 1, 1)
\overline{C}_1	(1, 0, 0)	(1, 1, 0)	\overline{C}_1	(1, 0, 1)	(1, 1, 1)

The explicit representation of this game in normal form would need exponential space, which illustrates the succinctness power of a representation by Boolean games.

Each player i has two possible strategies: $s_{i_1} = \{C_i\}$, $s_{i_2} = \{\overline{C}_i\}$. There are 8 strategy profiles for G . Consider $S_1 = (C_1, C_2, C_3)$ and $S_2 = (\overline{C}_1, C_2, C_3)$. Under S_1 , players 1, 2 and 3 have their goal satisfied, while S_2 satisfies only Player 1's goal.

This choice of binary utilities (where agents can only express plain satisfaction or plain dissatisfaction, with no intermediate levels) is a real loss of generality. We would like now to allow for associating an arbitrary preference relation on Ω with each player.

A preference relation \succeq is a reflexive and transitive binary relation (not necessarily complete) on Ω . The strict preference \succ associated with \succeq is defined as usual by $S_1 \succ S_2$ if and only if $S_1 \succeq S_2$ and not $(S_2 \succeq S_1)$.

A generalized Boolean game will be a 4-uple $G = (A, V, \pi, \Phi)$, where $A = \{1, \dots, n\}$, V and π are as before and $\Phi = \langle \Phi_1, \dots, \Phi_n \rangle$, where for each i , Φ_i is a compact representation (in some preference representation language) of the preference relation \succeq_i of agent i on Ω . We let $Pref_G = \langle \succeq_1, \dots, \succeq_n \rangle$.

A pure strategy Nash equilibrium (PNE) is a strategy profile such that each player's strategy is an optimum response to the other players' strategies. However, PNEs are classically defined for games where preferences are complete, which is not necessarily the case here. Therefore we have to define two notions of PNEs, a weak one and a strong one (they are equivalent to the notion of maximal and maximum equilibria in [14]).

Definition 3 Let $G = (A, V, \pi, \Phi)$ and $Pref_G = \langle \succeq_1, \dots, \succeq_n \rangle$ the collection of preference relations on Ω induced from Φ . Let $S = (s_1, \dots, s_n) \in \Omega$.

S is a **weak PNE** (WPNE) for G iff $\forall i \in \{1, \dots, n\}, \forall s'_i \in 2^{\pi_i}, (s'_i, s_{-i}) \not\succeq_i (s_i, s_{-i})$.

S is a **strong PNE** (SPNE) for G iff $\forall i \in \{1, \dots, n\}, \forall s'_i \in 2^{\pi_i}, (s'_i, s_{-i}) \preceq_i (s_i, s_{-i})$.

$NE_{strong}(G)$ and $NE_{weak}(G)$ denote respectively the set of strong and weak PNEs for G .

Clearly, any SPNE is a WPNE, that is, $NE_{strong}(G) \subseteq NE_{weak}(G)$.

3 Boolean games and prioritized goals

The preferences of a single player in this framework are expressed by a set of goals ordered by a priority relation:

Definition 4 A **prioritized goal base** Σ is a collection $\langle \Sigma^1; \dots; \Sigma^p \rangle$ of sets of propositional formulas. Σ^j represents the set of goals of priority j , with the convention that the smaller j , the more prioritary the formulas in Σ^j .

In this context, several criteria can be used in order to generate a preference relation \succeq from Σ . We recall below the three most common ones. In the following, if S is an interpretation of 2^V then we let $Sat(S, \Sigma^j) = \{\phi \in \Sigma^j \mid S \models \phi\}$.

Definition 5 Let $\Sigma = \langle \Sigma^1; \dots; \Sigma^p \rangle$, and let S and S' be two interpretations of 2^V .

Discrimin preference relation [7, 13, 2] $S \succ^{disc} S'$ iff $\exists k \in \{1, \dots, p\}$ such that:

$$Sat(S, \Sigma^k) \supset Sat(S', \Sigma^k) \text{ and } \forall j < k, Sat(S, \Sigma^j) = Sat(S', \Sigma^j)$$

Leximin preference relation [10, 2, 17] $S \succ^{lex} S'$ iff $\exists k \in \{1, \dots, p\}$ such that:

$$|Sat(S, \Sigma^k)| > |Sat(S', \Sigma^k)| \text{ and } \forall j < k, |Sat(S, \Sigma^j)| = |Sat(S', \Sigma^j)|.$$

Best-out preference relation [10, 2] Let $a(s) = \min\{j \text{ such that } \exists \varphi \in \Sigma^j, S \not\models \varphi\}$, with the convention $\min(\emptyset) = +\infty$. Then $S \succeq^{bo} S'$ iff $a(S) \geq a(S')$.

Note that \succeq^{bo} and \succeq^{lex} are complete preference relations, while \succeq^{disc} is generally a partial preference relation. Moreover, the following implications hold (see [2]):

$$(S \succ^{bo} S') \Rightarrow (S \succ^{disc} S') \Rightarrow (S \succ^{lex} S') \quad (1) \quad (S \succeq^{disc} S') \Rightarrow (S \succeq^{lex} S') \Rightarrow (S \succeq^{bo} S') \quad (2)$$

Definition 6 A **PG-Boolean game** is a 4-uple $G = (A, V, \pi, \Phi)$, where $\Phi = (\Sigma_1, \dots, \Sigma_n)$ is a collection of prioritized goals bases. We denote $\Sigma_i = \langle \Sigma_i^1, \dots, \Sigma_i^p \rangle$, that is, Σ_i^j denotes the stratum j of Σ_i , or equivalently, the (multi)set of goals of priority j for player i .


Note that the assumption that the number of priority levels is the same (p) for all players does not imply a loss of generality, as adding empty strata to a prioritized base does not change the induced preference relation.

We make use of the following notations:

- if G is a PG-boolean game and $c \in \{disc, lex, bo\}$ then $Pref_G^c = \langle \succeq_1^c, \dots, \succeq_n^c \rangle$.
- $NE_{weak}^c(G)$ and $NE_{strong}^c(G)$ denote respectively the sets of all weak and strong Nash equilibria for $Pref_G^c$.

Example 2 Let $G = (A, V, \pi, \Phi)$ with $A = \{1, 2\}$, $V = \{a, b, c\}$, $\pi_1 = \{a, c\}$, $\pi_2 = \{b\}$, $\Sigma_1 = \langle a; (\neg b, c) \rangle$, $\Sigma_2 = \langle (\neg b, \neg c); \neg a \rangle$.

For each of the three criteria $c \in \{lex, disc, bo\}$, we draw the corresponding preference relations $Pref_G^c = \langle \succeq_1^c, \dots, \succeq_n^c \rangle$. The arrows are oriented from more preferred to less preferred strategy profiles (S_1 is preferred to S_2 is denoted by $S_1 \rightarrow S_2$). To make the figures clearer, we do not draw edges that are obtained from others by transitivity. The dotted arrows indicate the links taken into account in order to compute Nash equilibria.



- **Discrimin and Leximin:** $NE_{weak}^{disc}(G) = NE_{strong}^{disc}(G) = \{\bar{a}bc\}$
- **Best Out:** $NE_{weak}^{bo}(G) = NE_{strong}^{bo}(G) = \{abc, \bar{a}bc\}$

Lemma 1 Let $\succeq = \langle \succeq_1, \dots, \succeq_n \rangle$ and $\succeq' = \langle \succeq'_1, \dots, \succeq'_n \rangle$ be two collections of preference relations, and let S be a strategy profile.

1. If \succeq is contained in \succeq' and if S is a SPNE for \succeq , then S is a SPNE for \succeq' .
2. If \succ is contained in \succ' and if S is a WPNE for \succ , then S is a WPNE for \succ' .

This lemma enables us to draw the following:

Proposition 1 Let $G = (A, V, \pi, \Phi)$ be a PG-boolean game and $Pref_G^c = \langle \succeq_1^c, \dots, \succeq_n^c \rangle$. $NE_{strong}^{disc}(G) \subseteq NE_{strong}^{lex}(G) \subseteq NE_{strong}^{bo}(G)$ and $NE_{weak}^{lex}(G) \subseteq NE_{weak}^{disc}(G) \subseteq NE_{weak}^{bo}(G)$.

We may now wonder whether a PG-boolean game can be *approximated* by focusing on the first k strata of each player. Here, the aim is double: to obtain a simpler (for PNE computation) game and to increase the possibility to find a significant PNE taking into account the most prioritized strata.

Definition 7 Let $G = (A = \{1, \dots, n\}, V, \pi, \Phi)$ be a PG-boolean game, and $k \in \{1, \dots, p\}$. $G^{[1 \rightarrow k]} = (A, V, \pi, \Phi^{[1 \rightarrow k]})$ denotes the k -reduced game of G in which all players' goals in G are reduced in their k first strata: $\Phi^{[1 \rightarrow k]} = \langle \Sigma_1^{[1 \rightarrow k]}, \dots, \Sigma_n^{[1 \rightarrow k]} \rangle$.

Lemma 2 Let G be a PG-boolean game. Then for every $k \leq p$, $c \in \{discr, lex, bo\}$, and every $i \in A$, we have: $S \succeq_i^{c, [1 \rightarrow k]} S' \Rightarrow S \succeq_i^{c, [1 \rightarrow k-1]} S'$ and $S \not\succeq_i^{c, [1 \rightarrow k]} S' \Rightarrow S \not\succeq_i^{c, [1 \rightarrow k-1]} S'$.

Proposition 2 Let G be a PG-boolean game and $c \in \{discr, lex, bo\}$. If S is a SPNE (resp. WPNE) for $Pref_G^c$ of the game $G^{[1 \rightarrow k]}$, then S is a SPNE (resp. WPNE) for $Pref_G^c$ of the game $G^{[1 \rightarrow (k-1)]}$.

This proposition leads in an obvious way that if $G^{[1]}$ for $Pref_G^c$ does not have any SPNE (resp. WPNE), then the game G for $Pref_G^c$ does not have any SPNE (resp. WPNE) whatever the criteria used. The converse is false, as shown in the following example.

Example 3 Let G with $A = \{1, 2\}$, $V = \{a, b\}$, $\pi_1 = \{a\}$, $\pi_2 = \{b\}$, $\Sigma_1 = \langle a \rightarrow b; b \rightarrow a \rangle$, $\Sigma_2 = \langle a \leftrightarrow \neg b; \neg b \rangle$. We check that $NE_{weak}^{bo}(G) = NE_{strong}^{bo}(G) = \emptyset$. Let us now focus on the 1-reduced game $G^{[1]} = (A, V, \pi, \Phi^{[1]})$ of G . We have $\Sigma_1^{[1]} = \langle a \rightarrow b \rangle$, $\Sigma_2^{[1]} = \langle a \leftrightarrow \neg b \rangle$. We check that for any criterion c , $NE_{weak}^c(G^{[1]}) = NE_{strong}^c(G^{[1]}) = \{\bar{a}b\}$.

This example shows us that Proposition 2 can be used to find the right level of approximation for a PG-game. For instance, we may want to focus on the largest k such that $G^{[1 \rightarrow k]}$ has a SPNE, and similarly for WPNEs.

4 Boolean games and CP-nets

A problem with prioritized goals is the difficulty for the agent to express his preferences (from a cognitive or linguistic point of view). In this Section we consider another very popular language for compact preference representation on combinatorial domains, namely CP-nets. This graphical model exploits conditional preferential independence in order to structure decision maker's preferences under a *ceteris paribus* assumption. They were introduced in [6] and extensively studied in many subsequent papers, especially [4, 5].

Although CP-nets generally consider variables with arbitrary finite domains, for the sake of simplicity (and homogeneity with the rest of the paper) here we consider only "propositionalized" CP-nets, that is, CP-nets with binary variables (note that this is not a real loss of generality, as all our definitions and results can be easily lifted to the more general case of non-binary variables).


Definition 8 Let V be a set of propositional variables and $\{X, Y, Z\}$ a partition of V . X is **conditionally preferentially independent** of Y given Z if and only if $\forall z \in 2^Z, \forall x_1, x_2 \in 2^X$ and $\forall y_1, y_2 \in 2^Y$ we have : $x_1 y_1 z \succeq x_2 y_1 z$ iff $x_1 y_2 z \succeq x_2 y_2 z$

For each variable X , the agent specifies a set of *parent variables* $\text{Pa}(X)$ that can affect her preferences over the values of X . Formally, X and $V \setminus (\{X\} \cup \text{Pa}(X))$ are conditionally preferentially independent given $\text{Pa}(X)$. This information is used to create the CP-net:

Definition 9 Let V be a set of variables. $\mathcal{N} = \langle \mathcal{G}, \mathcal{T} \rangle$ is a **CP-net on V** , where \mathcal{G} is a directed graph over V , and \mathcal{T} is a set of conditional preference tables $\text{CPT}(X_j)$ for each $X_j \in V$. Each $\text{CPT}(X_j)$ associates a total order \succ_p^j with each instantiation $p \in 2^{\text{Pa}(X_j)}$.

Definition 10 A **CP-boolean game** is a 4-uple $G = (A, V, \pi, \Phi)$, where $A = \{1, \dots, n\}$ is a set of players, $V = \{x_1, \dots, x_p\}$ is a set of variables and $\Phi = \langle \mathcal{N}_1, \dots, \mathcal{N}_n \rangle$. Each \mathcal{N}_i is a CP-net on V .

Example 4 $G = (A, V, \pi, \Phi)$ where $A = \{1, 2\}$ $V = \{a, b, c\}$ $\pi_1 = \{a, b\}$, $\pi_2 = \{c\}$, \mathcal{N}_1 and \mathcal{N}_2 are represented on the following figure.



Using these partial pre-orders, Nash equilibria are: $NE_{strong} = NE_{weak} = \{abc\}$.

The first property concerns a very interesting case where the existence and the unicity of PNE hold:

Proposition 3 Let $G = (A, V, \pi, \Phi)$ be a CP-boolean game such the graphs \mathcal{G}_i are all identical ($\forall i, j, \mathcal{G}_i = \mathcal{G}_j$) and acyclic. Then G has one and only one strong PNE.

The proof of this result makes use of the *forward sweep* procedure [6, 4] for outcome optimization (this procedure consists in instantiating variables following an order compatible with the graph, choosing for each variable its preferred value given the value of the parents).

The point is that in general the graphs \mathcal{G}_i for $i \in \{1, \dots, n\}$ may not be identical. However, they may be *made* identical, once remarked that a CP-net $\langle \mathcal{G}, \mathcal{T} \rangle$ can be expressed as a CP-net $\langle \mathcal{G}', \mathcal{T}' \rangle$ as soon as the set of edges in \mathcal{G} is contained in the set of edges in \mathcal{G}' . We may then take as common graph \mathcal{G} (to all players) the graph whose set of edges is the *union* of the set of edges of $\mathcal{G}_1, \dots, \mathcal{G}_n$. The only problem is that the resulting graph may not be acyclic, in which case Proposition 3 is not applicable. Formally:

Definition 11 Let G be a CP-boolean game. For each player i , \mathcal{G}_i is denoted by (V, Arc_i) , with Arc_i being the set of edges of i 's CP-net. The **union graph** of G is defined by $\mathcal{G} = (V, \text{Arc}_1 \cup \dots \cup \text{Arc}_n)$. The **normalized game equivalent to G** , denoted by $G^* = \{A, V, \pi, \Phi^*\}$, is the game obtained from

G by rewriting, where the graph of each player's CP-net has been replaced by the graph of the union of CP-nets of G and the CPT of each player's CP-net are modified in order to fit with the new graph, keeping the same preferences (formally, if \succ_i^y denotes the relation associated with $CPT_i(y)$ for Player i 's CP-net in G , then we have for G^* : $\forall x \in V$ such that x is a parent of y in G^* but not in G , $\succ_{i,x}^y = \succ_{i,\bar{x}}^y = \succ_i^y$).

The following lemma is straightforward:

Lemma 3 *Let G be a CP-boolean game and G^* its equivalent normalized game. Then G^* and G define the same preference relations on strategy profiles.*

Therefore, if G^* is acyclic, then Proposition 3 applies, therefore G^* has one and only one SPNE. Now, since G and G^* define the same pre-orders on Ω , the latter is also the only SPNE of G (on the other hand, if the graph of G is cyclic, neither the unicity nor the existence of SPNEs is guaranteed).

Proposition 4 *Let $G = (A, V, \pi, \Phi)$ be a CP-boolean game. If the union graph of G is acyclic then G has one and only one strong PNE.*

Example 4, continued: *Players' preferences in the normalized game G^* (equivalent to G) are represented by the CP-nets given on Figure 1. The union graph is acyclic, therefore Proposition 3 can be applied and G has one and only one strong PNE (abc).*

There is a last condition (less interesting in practice because it is quite strong) guaranteeing the existence and the unicity of a SPNE. This condition states that any variable controlled by an agent is preferentially independent on variables controlled by other agents (in other words, the parents of any variable controlled by a player i are also controlled by i). In this case, each agent is able to instantiate her variables in an unambiguously optimal way, according to her preferences.

Proposition 5 *Let $G = (A, V, \pi, \Phi)$ be a CP-boolean game such that for every player $i \in A$ and for every $v \in \pi_i$, we have $Pa(v) \in \pi_i$. Then G has one and only one SPNE.*

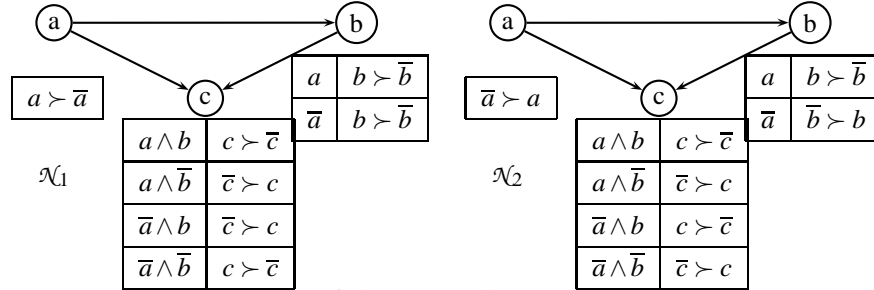


Figure 1: CP-net of Players 1 and 2's preferences for G^*

5 Related work and conclusion

Apart of previous work on Boolean games [15, 14, 11], related work includes a few papers where games are expressed within ordinal preferences within well-developed AI frameworks.

In [12], a game in normal form is mapped into a *logic program with ordered disjunction* (LPOD) where each player owns a set of clauses that encode the player's preference over her possible actions given every possible strategy profile of other players. It is shown that PNE correspond exactly to the most preferred answer sets. The given translation suffers from a limitation, namely its size: the size of the LPOD is the same as that of the normal form of the game (since each player needs a number of clauses equal to the number of possible other strategy profiles for other players). However, this limitation is due to the way LPODs are induced from games and could be overwhelmed by allowing to express the players' preferences by any LPODs (in the same spirit as our Section 3).

In [9], a strategic game is represented using a *choice logic program*, where a set of rules express that a player will select a “best response” given the other players’ choices. Then, for every strategic game, there exists a choice logic program such that the set of stable models of the program coincides with the set of Nash equilibria of the game. This property provides a systematic method to compute Nash equilibria for finite strategic games.

In [1], CP-nets are viewed as games in normal form and vice versa. Each player i corresponds to a variable X_i of the CP-net, whose domain $D(X_i)$ is the set of available actions to the player. Preferences over a player’s actions given the other players’ strategies are then expressed in a conditional preference table. The CP-net expression of the game can sometimes be more compact than its normal form explicit representation, provided that some players’ preferences depend only on the actions of a subset of other players. A first important difference with our framework is that we allow players to control an arbitrary set of variables, and thus we do not view players as variables; the only way of expressing in a CP-net that a player controls several variables would consist in introducing a new variable whose domain would be the set of all combination of values for these variables—and the size of the CP-net would then be exponential in the number of variables. A second important difference, which holds as well for the comparison with [12] and [9], is that players can express arbitrary preferences, including extreme cases where the satisfaction of a player’s goal may depend only of variables controlled by other players. A last (less technical and more foundational) difference with both lines of work, which actually explains the first two above, is that we do not *map* normal form games into anything but we *express* games using a logical language.

Further work includes the investigation of other notions (such as dominated strategies) within the two frameworks proposed in this paper, as well as the integration of other preference representation languages within Boolean games.

References

- [1] K. R. Apt, F. Rossi, and K. B. Venable. CP-nets and Nash equilibria. In Elsevier, editor, *Proc. CIRAS 2005 (Third International Conference on Computational Intelligence, Robotics and Autonomous Systems)*, Singapore, December 13-16 2005.
- [2] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proc. of the 13th IJCAI*, pages 640–645, 1993.
- [3] E. Bonzon, MC. Lagasquie-Schiex, J. Lang, and B. Zanuttini. Boolean games revisited. available at <ftp:ftp.irit.fr/pub/IRIT/RPDM/eca106.ps.gz>, 2006.
- [4] C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, and D. Poole. Cp-nets: A tool for representing and reasoning with conditional *Ceteris Paribus* preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.
- [5] C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, and D. Poole. Preference-based constrained optimization with cp-nets. *Computational Intelligence*, 20(2):137–157, 2004. Special Issue on Preferences.
- [6] C. Boutilier, R. I. Brafman, H. H. Hoos, and D. Poole. Reasoning with conditional *ceteris paribus* preference statements. In *Proc. of UAI*, 1999.
- [7] G. Brewka. Preferred subtheorie : An extended logical framework for default reasoning. In *Proc. of the 11th IJCAI*, pages 1043–1048, 1989.
- [8] S. Coste-Marquis, J. Lang, P. Liberatore, and P. Marquis. Expressive power and succinctness of propositional languages for preference representation. In *Proc. of the 9th KR*, pages 203–212, 2004.
- [9] M. De Vos and D. Vermeir. Choice logic programs and Nash equilibria in strategic games. In Jorg Flum and Mario Rodriguez-Artalejo, editors, *Computer Science Logic (CSL’99)*, volume 1683, pages 266–276, 1999.
- [10] D. Dubois, J. Lang, and H. Prade. Inconsistency in possibilistic knowledge bases: To live with it or not live with it. In *Fuzzy Logic for the Management of Uncertainty*, pages 335–351. Wiley and sons, 1992.
- [11] P.E. Dunne and W. van der Hoek. Representation and complexity in boolean games. In *Proc. of JELIA, LNCS 3229*, pages 347–359, 2004.
- [12] N. Foo, T. Meyer, and G. Brewka. LPOD answer sets and Nash equilibria. In M. Maher, editor, *Proceedings of the 9th Asian Computer Science Conference (ASIAN 2004)*, pages 343–351. Chiang Mai, Thailand, Springer LNCS 3321, 2004.

- [13] H. Geffner. Default reasoning: causal and conditional theories. *MIT Press*, 1992.
- [14] P. Harrenstein. *Logic in Conflict*. PhD thesis, Utrecht University, 2004.
- [15] P. Harrenstein, W. van der Hoek, J.J Meyer, and C. Witteveen. Boolean games. In *Proc. of TARK*, pages 287–298, 2001.
- [16] J. Lang. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42:37–71, 2004.
- [17] D Lehmann. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, 15:61–82, 1995.